

Efficient Circuits for Exact-Universal Computation With Qudits

Gavin K. Brennen¹, Stephen S. Bullock^{2,3} and Dianne P. O’Leary^{2,4}

September 22, 2005

¹Atomic Physics Division, National Institute of Standards and Technology, Gaithersburg, MD 20899-8420

²Mathematical and Computational Sciences Division, National Institute of Standards and Technology
Gaithersburg, MD 20899-8910

³Center for Computing Sciences, Institute for Defense Analyses, Bowie, MD 20715-4300

⁴Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742.

Abstract

This paper concerns the efficient implementation of quantum circuits for qudits. We show that controlled two-qudit gates can be implemented without ancillas and prove that the gate library containing arbitrary local unitaries and one two-qudit gate, CINC, is exact-universal. A recent paper [S.Bullock, D.O’Leary, and G.K. Brennen, Phys. Rev. Lett. **94**, 230502 (2005)] describes quantum circuits for qudits which require $O(d^n)$ two-qudit gates for state synthesis and $O(d^{2n})$ two-qudit gates for unitary synthesis, matching the respective lower bound complexities. In this work, we present the state synthesis circuit in much greater detail and prove that it is correct. Also, the $\lceil (n-2)/(d-2) \rceil$ ancillas required in the original algorithm may be removed without changing the asymptotics. Further, we present a new algorithm for unitary synthesis, inspired by the QR matrix decomposition, which is also asymptotically optimal.

1 Introduction

A qudit is a d -level generalization of a qubit, i.e. the one-qudit Hilbert space splits orthogonally as

$$\mathcal{H}(1, d) = \mathbb{C}\{|0\rangle\} \oplus \mathbb{C}\{|1\rangle\} \oplus \cdots \oplus \mathbb{C}\{|d-1\rangle\} \quad (1)$$

while the n -qudit state-space is $\mathcal{H}(n, d) = [\mathcal{H}(1, d)]^{\otimes n}$. Thus for $N = d^n$, closed-system evolutions of n qudits are modeled by $N \times N$ unitary matrices. Qudit circuit diagrams then factor such unitaries into two-qudit operations $I_{d^{n-2}} \otimes V$ where V is a $d^2 \times d^2$ unitary matrix, or more generally into similarity transforms of such gates by particle-swaps. The algorithmic complexity of an evolution may then be thought of as the number of two-qudit gates required to build it. A degree of freedom argument [9] leads one to guess that exponentially many gates are required for most unitary evolutions, since the space of all $N \times N$ unitary matrices is d^{2n} -dimensional. Indeed, this space of evolutions is a manifold so the argument may be made rigorous using smooth topology, and thus $\Omega(d^{2n})$ gates are required for exact-universality. Yet until quite recently the best qudit circuits contained $O(n^2 d^{2n})$ gates [10]. In contrast, $O(4^n)$ gates were known to suffice for qubits ($d = 2$) [14], presenting the possibility that qudits are genuinely less efficient for d not a power of two.

Quite recently, an explicit $O(d^{2n})$ construction was achieved [4]. It uses the spectral decomposition of the unitary matrix desired and also a new *state synthesis circuit* [6, 9, 11, 2]. Given a $|\psi\rangle \in \mathcal{H}(n, d)$, a state-synthesis circuit for $|\psi\rangle$ realizes some unitary U such that $U|\psi\rangle = |0\rangle$. There are $2d^n - 2$ real degrees of freedom in a normalized state ket $|\psi\rangle$, which may be used to prove that circuits for generic states cost $\Omega(d^n)$ two-qudit gates. This is in sharp contrast to the case of classical logic, where $O(n)$ inverters may produce any bit-string. The most recent qudit state-synthesis circuit [4] contains $(d^n - 1)/(d - 1)$ two-qudit gates, and in fact each is a singly-controlled one-qudit operator $\wedge_1(V) = I_{d^2-d} \oplus V$.

There are two ways to employ an asymptotically optimal state synthesis circuit in order to obtain asymptotically optimal unitary circuits. The first is to exploit the spectral decomposition, which involves a three part circuit for each eigenstate of the unitary: building an eigenstate [9, 4], applying a conditional phase to one logical basis ket, and unbuilding the eigenstate. We here introduce a second option, the **Triangle** algorithm, which uses the state-synth circuit with extra controls to reduce the unitary to upper triangular form. Recursive counts of the number of control boxes show that it is also asymptotically optimal (Cf. [14].) Although these algorithms are unlikely to be used to implement general unitary matrices, they can be usefully applied to improving subblocks of larger circuits (peephole optimization).

Finally, this work also addresses two further topics in which qudit circuits lag behind qubit circuits. First, to date the smallest gate library for exact universality with qudits uses arbitrary locals complemented by a continuous one parameter two-qudit gate [3]. In contrast, it is well known [6] that any computation on qubits can be realized using gates from the library $\{U(2)^{\otimes n}, \text{CNOT}\}$. We prove that the library $\{U(d)^{\otimes n}, \text{CINC}\}$, where CINC is the qudit generalization of the CNOT gate, is exactly universal. Second, the first asymptotically optimal qubit quantum circuit exploited a single ancilla qubit [14] and current constructions require none [2, 11], while qudit diagrams tend to suppose $\lceil (n-2)/(d-2) \rceil$ ancilla qudits. Here we present methods which realize a k -controlled operation $\wedge_k(V) = I_{d^{k+1}-d} \oplus V$ in $O[(k+2)^{2+\log_2 d}]$ gates without the need for any ancilla. This makes all qudit asymptotics competitive with their qubit counterparts. However, it is not known whether the Cosine-Sine Decomposition (CSD) is useful for building qudit circuits, despite the fact that all best-practice qubit exact universal circuits exploit this matrix decomposition.

The paper is organized as follows. §3 improves on earlier constructions of $\wedge_1(V)$ gates, which are ubiquitous in later sections. §4 presents a new circuit for a qudit $\wedge_k(V)$ gate which are later used to produce the first $O(d^{2n})$ gate unitary circuits without ancilla. §5 details the recent state synthesis algorithm as an iteration over a new \clubsuit -sequence and exploits the new constructions to prove it is correct. §6 presents a new asymptotically optimal unitary circuit inspired by the QR matrix factorization and compares it with a previous algorithm based on spectral decomposition. §7 discusses two applications of the state synthesis algorithm

2 Notation and conventions

The Hilbert spaces $\mathcal{H}(1, d)$ and $\mathcal{H}(n, d)$ are defined in the introduction. On $\mathcal{H}(1, d)$, the inverter for bits has two important generalizations for dits:

$$\begin{aligned} \sigma_x \oplus I_{d-2} |j\rangle &= \begin{cases} |1\rangle & j=0 \\ |0\rangle & j=1 \\ |j\rangle & 2 \leq j \leq d-1 \end{cases} \\ \text{INC} |j\rangle &= |(j+1) \bmod d\rangle \end{aligned} \tag{2}$$

We use the latter symbol rather than the more typical X since this operation is a modular increment. This leads to two generalizations of the quantum controlled-not, $\wedge_1(\sigma_x \oplus I_{d-2})$ and $\wedge_1(\text{INC}) = \text{CINC}$. The usual symbol for a controlled-not when appearing in a qudit circuit diagram refers to CINC. Controls represented by a black bubble in qudit circuit diagrams fire on control state $|d-1\rangle$.

As new notation, the \clubsuit -sequence is introduced in §5.2. This plays a role analogous to the Gray code in earlier $d=2$ constructions and is a particular sequence of words of n -letters. Although these words might themselves be called sequences, we prefer to call an individual word (e.g. 1100 $\clubsuit\clubsuit$) a *term* and reserve “sequence” exclusively for the \clubsuit -sequence of $(d^n - 1)/(d - 1)$ terms.

3 Optimizing singly-controlled one-qudit unitaries

Several operators $\wedge_1(V)$ appear in later circuits. Thus, it is worthwhile to optimize this computation in our gate libraries. For qubits, CNOT-optimal circuits for $\wedge_1(V)$ are known [12]. The qudit case is open. Here we improve the $\wedge_1(V)$ circuit in that work and further prove for the first time that $U(d)^{\otimes n} \sqcup \{\text{CINC}, \text{CINC}^{-1}\}$ is exact-universal.

Since $\text{CINC}^{-1} = \text{CINC}^{d-1}$, this also demonstrates that $U(d)^{\otimes n} \sqcup \{\text{CINC}\}$ is exact-universal. This is a smaller universal library than that presented in earlier work [3].

Thus, consider the question of factoring $\wedge_1(V)$. Let $\{|\psi_k\rangle\}_{k=0}^{d-1}$ be the eigenkets of V with eigenvalues $\{e^{i\theta_k}\}_{k=0}^{d-1}$. Let W_k be some one-qudit unitary with $W_k|0\rangle = |\psi_k\rangle$, e.g. the appropriate one-qudit Householder reflection (See §5.1.) Finally, let Φ_k be a controlled one-qudit phase unitary given by $\Phi_k = \wedge_1[I_d + (e^{i\theta_k} - 1)|0\rangle\langle 0|]$. Then note that $V = \prod_{k=0}^{d-1} W_k[I_d + (e^{i\theta_k} - 1)|0\rangle\langle 0|]W_k^\dagger$. Thus $\wedge_1(V)$ can be implemented by the following circuit:

$$\text{Circuit for } \wedge_1(V) \cong W_0^\dagger \Phi_0 W_0 W_1^\dagger \Phi_1 W_1 \cdots W_{d-1}^\dagger \Phi_{d-1} W_{d-1} \quad (3)$$

Thus, we have reduced the question to building Φ_k in terms of $U(d)^{\otimes n}$ and CINC .

Building Φ_k requires some preliminary remarks. Suppose we have $\xi \in \mathbb{C}$, $|\xi| = 1$. Consider the diagonal unitary of the corresponding geometric sequence: $D = \sum_{j=0}^{d-1} \xi^j |j\rangle\langle j|$. Recall that INC is the increment permutation, i.e. $\text{INC}|j\rangle = |(j+1) \bmod d\rangle$. Thus permuting the diagonal entries, $\text{INC} D \text{INC}^{-1} = \xi^{d-1}|0\rangle\langle 0| + \sum_{j=1}^{d-1} \xi^{j-1}|j\rangle\langle j|$. Hence

$$\text{INC} D \text{INC}^{-1} D^{-1} = \xi^{d-1}|0\rangle\langle 0| + \xi^{-1} \sum_{j=1}^{d-1} |j\rangle\langle j| = (\xi^{-1} I_d) (\xi^d |0\rangle\langle 0| + \sum_{j=1}^{d-1} |j\rangle\langle j|). \quad (4)$$

Now generalizing a standard trick from qubits, note further that

$$\wedge_1(\xi I_d) = \left(\sum_{j=0}^{d-2} |j\rangle\langle j| + \xi |d-1\rangle\langle d-1| \right) \otimes I_d, \quad (5)$$

so that a controlled global-phase is in fact a local operation. Hence taking $\xi = e^{i\theta_k/d}$, we obtain in particular an expression for Φ_k of Equation 3 in terms of CINC and CINC^{-1} :

$$\begin{aligned} \Phi_k &= \wedge_1(\xi I_d) \text{CINC} (I_d \otimes D) \text{CINC}^{-1} (I_d \otimes D^{-1}) \\ &= \left[\left(\sum_{j=0}^{d-2} |j\rangle\langle j| + \xi |d-1\rangle\langle d-1| \right) \otimes I_d \right] \text{CINC} (I_d \otimes D) \text{CINC}^{-1} (I_d \otimes D^{-1}). \end{aligned} \quad (6)$$

Hence, $\wedge_1(V)$ may be realized using gates from $U(d)^{\otimes n}$ along with d copies of CINC and d copies of CINC^{-1} .

Recall that these circuits may be expanded into circuits in terms of $\wedge_1(\sigma_x \oplus I_{d-2})$. Indeed, when viewed as permutations, INC and INC^{-1} factor into d flips. To see this, consider $0 \leq j < k \leq d-1$ and let (jk) denote the flip permutation $j \leftrightarrow k$ of $\{0, 1, \dots, d-1\}$. Then

$$\text{INC} = (01) \circ (12) \circ \cdots \circ (d-2 \ d-1). \quad (7)$$

Since $\wedge_1[(jk)]$ is equivalent to $\wedge_1(\sigma_x \oplus I_{d-2})$ up to permutations within $U(d)^{\otimes n}$, we see that CINC and CINC^{-1} may be implemented using $d-1$ copies of the controlled-flip. Thus, $\wedge_1(V)$ may also be realized using $2d(d-1)$ copies of the $\wedge_1(\sigma_x \oplus I_{d-2})$ gate.

Remark: Note that the controlled-flip is also equivalent to $\wedge_1(I_{d-2} \oplus \sigma_z)$, making blockwise use of the 2×2 matrix identity $H\sigma_x H = \sigma_z$ for $H = \frac{1}{\sqrt{2}} \sum_{j,k=0}^1 (-1)^{jk} |k\rangle\langle j|$. Thus, the above also realizes $\wedge_1(V)$ in roughly $2d^2$ controlled- π phase gates. This is half the roughly $4d^2$ gates of earlier work [3], even after including the arbitrary relative phase $e^{i\theta}$ allowed there.


4 Qudit control without ancillas

In this section we simulate a $\wedge_{n-1}(V)$ gate for $V \in U(d)$ using $O[(n+1)^{\log_2 d+2}]$ singly-controlled one qudit gates without ancilla. The method parallels the techniques used in Ref. [1] for universal computation with qubits.

First we decompose a $\wedge_{n-1}(V)$ gate using a sequence of gates with a smaller number of controls. As a first step, notice that

$$\wedge_{n-1}(V) = \wedge_{n-2}(X_{n-1}) [\wedge_{n-2}(\text{INC}) \wedge_1(X_{n-1}^\dagger)]^{d-1} \wedge_{n-2}(\text{INC}) \wedge_1(X_{n-1}^{d-1}), \quad (8)$$


where $X_{n-1} = V^{1/d}$. For example, for $n = 7$, we have the following circuit:



All control operations are conditioned on the control qudits being in state $|d-1\rangle$. The circuit is designed to cycle over each possible dit value of the control qudit in the $\wedge_1(X_{n-1})$ gates. The entire construction then follows by recursive application of Equation 8 to the last gate. In theory, this construction is an exact implementation of $\wedge_{n-1}(V)$. Yet in practice, the sequence of matrices X_j obtained by taking the d -th root of X_{j+1} (with $X_n = V$) quickly converges to the identity matrix as j decreases. Hence, an approximate implementation results if the recursion is terminated early.

As an example of Equation 9, consider the generalized Toffoli gate $\wedge_2(\text{INC})$. This breaks into $(d+1)$ variants of singly-controlled $\wedge_1(W)$ gates along with d extra CINC gates. Hence $(d+1)d + d$ CINC gates along with $(d+1)d$ CINC $^{-1}$ gates and sundry gates from $U(d)^{\otimes n}$ suffice to emulate $\wedge_2(\text{INC})$.

Note that the size of the circuit for $\wedge_{n-2}(\text{INC})$ that is analogous to the above grows exponentially in n . However, it is possible to simulate $\wedge_{n-2}(\text{INC})$ more efficiently using a sequence of $\wedge_{\lceil (n-1)/2 \rceil}(\text{INC})$ and $\wedge_{\lfloor (n-1)/2 \rfloor}(\text{INC})$ gates, proceeding recursively down to $\wedge_2(\text{INC})$. The argument is analogous to that used for qubits in Lemma 7.3 in Ref. [1] for $n \geq 5$. The following circuit illustrates the method for $n = 7$:



Ignoring which qudits are controlled or targeted, the circuit sequence is $\wedge_{n-2}(\text{INC}) = [\wedge_{\lfloor (n-1)/2 \rfloor}(\text{INC}) \wedge_{\lceil (n-1)/2 \rceil}(\text{INC})]^d$.

For the remainder of this section, we use a tilde to distinguish a count for CINC $^{-1}$ from a CINC count. Thus, we let b_{n-2} be the total number of CINC gates required to emulate $\wedge_{n-2}(\text{INC})$, and \tilde{b}_{n-2} be the similar count for CINC $^{-1}$. For Circuit 10,

$$\begin{aligned} b_{n-2} &= d(b_{\lceil (n-1)/2 \rceil} + b_{\lfloor (n-1)/2 \rfloor}), \\ \tilde{b}_{n-2} &= d(\tilde{b}_{\lceil (n-1)/2 \rceil} + \tilde{b}_{\lfloor (n-1)/2 \rfloor}). \end{aligned} \quad (11)$$

A quick induction shows that each sequence is increasing, and thus $b_{n-2} \leq 2db_{\lceil (n-1)/2 \rceil}$ and $\tilde{b}_{n-2} \leq 2d\tilde{b}_{\lceil (n-1)/2 \rceil}$. Moreover, by the analysis of $\wedge_2(\text{INC})$ above $b_2 = d^2 + 2d$ and $\tilde{b}_2 = d^2 + d$. Recalling $(\log_d n)(\log_2 d) = \log_2 n$, we obtain the following:

$$\begin{aligned} b_{n-2} &\leq (d^2 + 2d)(2d)(2d)^{\log_2 n} = (d^2 + 2d)(2d)n^{1+\log_2 d}, \\ \tilde{b}_{n-2} &\leq (d^2 + d)(2d)(2d)^{\log_2 n} = (d^2 + d)(2d)n^{1+\log_2 d}. \end{aligned} \quad (12)$$

Note that these counts assume that the emulation of $\wedge_{n-2}(\text{INC})$ is done on a system with n qudits. Combining this circuit with Circuit 9 allows for an ancilla-free implementation of $\wedge_{n-1}(V)$.

Thus, let c_{n-1} be the number of CINC gates required to emulate $\wedge_{n-1}(V)$, not counting an additional \tilde{c}_{n-1} CINC $^{-1}$ gates. Using Circuit 9,

$$\begin{aligned} c_{n-1} &= db_{n-2} + c_{n-2} + d^2, \\ \tilde{c}_{n-1} &= d\tilde{b}_{n-2} + \tilde{c}_{n-2} + d^2. \end{aligned} \quad (13)$$

We may then overestimate c_{n-1} and \tilde{c}_{n-1} using integral comparison and $c_2 = d^2 + 2d$, $\tilde{c}_2 = d^2 + d$, obtaining

$$\begin{aligned} c_{n-1} &= d\left(\sum_{j=2}^{n-2} b_j\right) + c_2 + (n-3)d^2 \\ &\leq d[(d^2 + 2d)(2d)] \int_4^{n+1} t^{1+\log_2 d} dt + 2d + (n-2)d^2 \\ &= \frac{(2d^2)(d^2+2d)}{2+\log_2 d} [(n+1)^{2+\log_2 d} - 4d^2] + (n-2)d^2 + 2d. \end{aligned} \quad (14)$$

We may similarly overestimate \tilde{c}_{n-1} :

$$\tilde{c}_{n-1} \leq \frac{(2d^2)(d^2+d)}{2+\log_2 d} [(n+1)^{2+\log_2 d} - 4d^2] + (n-2)d^2 + d. \quad (15)$$

Hence c_{n-1} , \tilde{c}_{n-1} are both bounded by $O[(n+1)^{2+\log_2 d}]$. This can be used to show that the earlier spectral algorithm [4] is asymptotically optimal even when ancilla qudits are absent.

If we disallow CINC^{-1} and rather emulate $\text{CINC}^{-1} = \text{CINC}^{d-1}$, then the overall CINC count for $\wedge_{n-1}(V)$ would be $c_{n-1} + (d-1)\tilde{c}_{n-1}$. Note that if the gate library contains the two qudit gate $\wedge_1(\sigma_x \oplus I_{d-2})$ rather than CINC , a naïve application of the above argument would imply a linear overhead with a factor of $d-1$. However Circuits 9 and 10 can be adapted by replacing the $\wedge_k(\text{INC})$ gates with gates locally equivalent to $\wedge_1(\sigma_x \oplus I_{d-2})$, resulting in a smaller overhead.

5 Asymptotically optimal qudit state synthesis

State-synthesis is an important problem in quantum circuit design [6, 9]. This section expands upon the earlier account [4] of an asymptotically optimal state synthesis circuit for qudits. The earlier circuit used only $O(d^n)$ two-qudit gates, while a dimension-based argument [4] shows that no fewer ($\Omega(d^n)$) gates may achieve qudit state synthesis. There are two extensions in the present account:

- We introduce the ♣-sequence, a combinatorial gadget that organizes the order in which amplitudes are zeroed while (de)constructing the target state.
- Using the ♣-sequence, we prove that the state synthesis algorithm functions as asserted.

The two-qudit gates are in fact all $\wedge_1(V)$ for V a one-qudit Householder reflection. Hence, earlier sections of the present work further improve the previous circuit.

Recall from the introduction that we prefer to build W with $W|\psi\rangle = |0\rangle$ rather than building U with $U|0\rangle = |\psi\rangle$. We do this by constructing a sequence of factors which introduce more zeros into the partially zeroed state. The ordering established here by the ♣-sequence may be replaced by Gray code ordering [14] in the case $d = 2$.

5.1 One-qudit Householder reflections

Earlier universal $d = 2$ circuits [1] relied on a QR factorization to write any unitary U as a product of *Givens rotations*, realized in the circuit as k -controlled unitaries [5]. In the multi-level case, we instead use *Householder reflections* [7, §5.1]. Thus, suppose $|\psi\rangle \in \mathcal{H}(1, d)$, perhaps not normalized. Householder reflections solve the one-qudit case of the inverse state-synthesis problem. Suppose

$$\begin{cases} |\eta\rangle &= |\psi\rangle - \sqrt{\langle\psi|\psi\rangle} \frac{\langle 0|\psi\rangle}{|\langle 0|\psi\rangle|} |0\rangle, \\ W &= I_d - (2/\langle\eta|\eta\rangle) |\eta\rangle\langle\eta|. \end{cases} \quad (16)$$

Then $W|\psi\rangle$ is a multiple of $|0\rangle$. Geometrically, W is that unitary matrix which reflects across a plane lying between $|0\rangle$ and $|\psi\rangle$.

5.2 Inserting zeroes using Householders in ♣-sequence order

The n -qudit techniques require a bit more notation. Any term of the ♣-sequence describes a particular instantiation of a $\wedge_k(V)$ gate, controlled on certain lines determined by the letters with target determined by the first ♣. We next expand the controlled operator notation so as to precisely describe how to extract a control from such a term.

Definition 5.1 [4] [Controlled one-qudit operator $\wedge(C, V)$] Let V be a $d \times d$ unitary matrix, i.e. a one-qudit operator. Let $C = [C_1 C_2 \dots C_n]$ be a length- n control word composed of letters from the alphabet $\{0, 1, \dots, d-1\} \sqcup \{*\} \sqcup \{T\}$, with exactly one letter in the word being T . By $\#C$ we mean the number of letters in the word with numeric values (i.e., the number of controls.) The set of control qudits is the corresponding subset of $\{1, 2, \dots, n\}$ denoting the positions of numeric values in the word. A control word matches an n -dit string if each numeric value matches. Then the controlled one-qudit operator $\wedge(C, V)$ is the n -qudit operator that applies V to the qudit specified by the position of T iff the control word matches the data state's n -dit string. More precisely, in the case when $C_n = T$, then

$$\wedge([C_1 C_2 \dots C_{n-1} T], V) |c_1 c_2 \dots c_n\rangle = \begin{cases} |c_1 \dots c_{n-1}\rangle \otimes V |c_n\rangle, & c_j = C_j \text{ or } C_j = *, 1 \leq j \leq n-1 \\ |c_1 \dots c_{n-1} c_n\rangle, & \text{otherwise} \end{cases} \quad (17)$$

Alternatively, if $C_j = T$ ($j < n$), we consider the unitary (permutation) operator χ_j^n that swaps qudits j and n . Thus, $\chi_j^n |d_1 d_2 \dots d_n\rangle = |d_1 d_2 \dots d_{j-1} d_n d_{j+1} \dots d_{n-1} d_j\rangle$. Control on a word $C = [C_1 C_2 \dots C_{j-1} T C_{j+1} \dots C_n]$, is then given by $\wedge(C, V) = \chi_j^n \wedge(\tilde{C}, V) \chi_j^n$ for $\tilde{C} = [C_1 C_2 \dots C_{j-1} C_n C_{j+1} \dots C_{n-1} T]$.

In our particular state synthesis algorithm, we can factor W so that $\prod_{k=1}^p \wedge[C(p-k+1), V(p-k+1)] |\psi\rangle = |0\rangle$ with all $\#C(k) \leq 1$ and $p = (d^n - 1)/(d - 1)$. Since each $\#C(k) \leq 1$, each controlled operation is in fact a two-qudit gate. The circuit layout depends on the ♣-sequence, defined in Algorithm 1 and illustrated in Table 1.

n	♣-sequence, $d = 3$
1	♣
2	0♣, 1♣, 2♣, ♣♣
3	00♣, 01♣, 02♣, 0♣♣, 10♣, 11♣, 12♣, 1♣♣, 20♣, 21♣, 22♣, 2♣♣, ♣♣♣
4	000♣, 001♣, 002♣, 00♣♣, 010♣, 011♣, 012♣, 01♣♣, 020♣, 021♣, 022♣, 02♣♣, 0♣♣♣, 100♣, 101♣, 102♣, 10♣♣, 110♣, 111♣, 112♣, 11♣♣, 120♣, 121♣, 122♣, 12♣♣, 1♣♣♣, 200♣, 201♣, 202♣, 20♣♣, 210♣, 211♣, 212♣, 21♣♣, 220♣, 221♣, 222♣, 22♣♣, 2♣♣♣, ♣♣♣♣

Table 1: Sample ♣-sequences for $d = 3$, i.e. qutrits.

Algorithm 1: $\{s_1, \dots, s_p\} = \mathbf{Make\text{-}\clubsuit\text{-sequence}}(d, n)$

% We return a sequence of $p = (d^n - 1)/(d - 1)$ terms, with n letters each,

% drawn from the alphabet $\{0, 1, \dots, d-1, \clubsuit\}$.

Let $\{\tilde{s}_j\}_{j=1}^{\tilde{p}} = \mathbf{Make\text{-}\clubsuit\text{-sequence}}(d, n-1)$.

for $q = 0, 1, \dots, d-1$ **do**

 The next $(d^{n-1} - 1)/(d - 1)$ terms of the sequence are formed by prefixing the letter q to each term of the sequence $\{\tilde{s}_j\}$.

end for

The final term of the sequence is ♣ ^{n} .

The number of elements in the sequence, $(d^n - 1)/(d - 1)$, equals the number of uncontrolled or singly-controlled one-qudit operators in our state-synthesis circuit. To produce the circuit, it suffices to describe how to extract the control word C from a term t of the ♣-sequence and how to determine V from the term and $|\psi_j\rangle$, where $|\psi_j\rangle = \prod_{k=1}^{j-1} \wedge[C(p-k+1), V(p-k+1)] |\psi\rangle$ is the partial product, as shown in the following algorithm.

Algorithm 2: $\wedge(C, V) = \text{Single-}\clubsuit\text{Householder}(\clubsuit \text{ term } t = t_1 t_2 \dots t_n, n\text{-qudit state } |\psi_j\rangle)$

```

Initialize  $C = * * \dots *$ 
% Set the target:
Let  $\ell$  be the index of the leftmost  $\clubsuit$  and set  $C_\ell = T$ .
% Set a single control if needed:
if  $t$  contains numeric values greater than 0,
    Let  $q$  be the index of the rightmost such value and set  $C_q = t_q$ .
end if
Given  $|\psi_j\rangle = \sum_{k=0}^{d^n-1} \langle k | \psi_j \rangle |k\rangle$ , form a one-qudit state  $|\phi\rangle = \sum_{k=0}^{d-1} \langle t_1 t_2 \dots t_{\ell-1} k 00 \dots 0 | \psi_j \rangle |k\rangle$ .
Form  $V$  as a one-qudit Householder such that  $V|\phi\rangle = |0\rangle$ .

```

Figure 1 displays the type of gate produced from the output C and V from the algorithm **Single- \clubsuit Householder**. Figure 2 illustrates the order in which these $\wedge(C, V)$ reflections are generated if we iterate over the \clubsuit -sequence. Each node of the tree is labeled by a \clubsuit -term and represents a Householder reflection defined by three elements of $|\psi\rangle$, whose indices are indicated in the node. The reflection zeroes all but the first of these three elements. The reflections are applied by traversing the graph in depth-first order, left to right.




Figure 1: Producing a $\wedge(C, V)$ given V and a term of the \clubsuit -sequence, here $t = 2100\clubsuit\clubsuit\clubsuit$ for seven qudits. The algorithm for producing C places the V -target symbol T on the leftmost club, here line 5. The active control must then be placed on the least significant line carrying a nonzero prior to line 5, here the 1 on line 2. (A control on lines 3 or 4 would not prevent the nonzero α_0 of $|\psi_j\rangle = \sum_{k=0}^{d^n-1} \alpha_k |k\rangle$ from creating new nonzero entries in previously zeroed positions.) Thus in this case, $C = *1**T**$. The V is chosen to zero all but one α_k for $k = 2100\ell 00$.

5.3 Householder circuits for state synthesis

We will make use of state synthesis for $|\psi\rangle \mapsto \sqrt{\langle\psi|\psi\rangle}|0\rangle$ but also for $|\psi\rangle \mapsto \sqrt{\langle\psi|\psi\rangle}|m\rangle$ for any $m = d_1 d_2 \dots d_n$. We adapt our construction for a collapse onto $|0\rangle$ into an algorithm for collapse onto $|m\rangle$. The idea is to permute the elements to put m in position 0, apply a **Single- \clubsuit Householder** sequence, and then permute back.

Let $m = d_1 d_2 d_3 \dots d_n$ be a d -ary expansion of some m , $0 \leq m \leq d^n - 1$. Then $|m\rangle = \otimes_{k=1}^n \text{INC}^{d_k} |0\rangle$. Further, for a generic control word C , define a new m -dependent control word \tilde{C} by

$$\tilde{C}_k = \begin{cases} *, & C_k = * \\ T, & C_k = T \\ (C_k + d_k) \bmod d, & C_k \in \{0, 1, \dots, d-1\} \end{cases} \quad (18)$$

Suppose also that $C_m = T$. Then noting that $(\oplus m)^\dagger = \oplus(d - m)$, we have the similarity relation

$$[\otimes_{k=1}^n \text{INC}^{d_k}] \wedge(C, V) [\otimes_{k=1}^n \text{INC}^{d-d_k}] = \wedge[\tilde{C}, (\oplus d_m) V (\oplus d - d_m)]. \quad (19)$$

This is the basis for the algorithm for state synthesis.




Figure 2: Using the \clubsuit -sequence for $d = 3$, $n = 3$ to generate Householder reflections to reduce $|\psi\rangle$ to a multiple of $|0\rangle$. Each node is labeled by a \clubsuit -term and represents a Householder reflection $\wedge(C, V)$. The control is indicated by the boldface entry in the label. As the tree is traversed in a depth-first search, each node indicates a $\wedge(C, V)$ that zeroes the components of the last two indices in each node using the component of the top entry. See also Figure 1 of [4].

Algorithm 3: $\wedge(C, V) = \clubsuit\text{Householder}(|\psi\rangle, m, d, n)$

% Reduce $|\psi\rangle$ onto $|m\rangle$.
 Let $m = d_1 d_2 \dots d_n$.
 Compute $|\phi\rangle = (\otimes_{q=1}^n \text{INC}^{d-d_q}) |\psi\rangle$.
 Produce a sequence of controlled one-qudit operators so that
 $\prod_{k=1}^p \wedge[C(p-k+1), V(p-k+1)] |\phi\rangle = |00\dots 0\rangle$,
 using **Single- \clubsuit Householder** applied to each term of **Make- \clubsuit -sequence**(d, n).
 Compute $(\otimes_{q=1}^n \text{INC}^{d_q}) \wedge[C(p-k+1), V(p-k+1)] (\otimes_{p=1}^n \text{INC}^{d-d_q}) =$
 $\wedge[\tilde{C}(p-k+1), \tilde{V}(p-k+1)]$ using Equation 19

\clubsuit Householder applies the sequence of Householder reflections generated by **Single- \clubsuit Householder**. The resulting unitary W , although not a Householder reflection itself, satisfies $W|\psi\rangle = |m\rangle$, as we prove in the next subsection. Moreover, since the circuit contains $O(d^n)$ two-qudit gates, all of which are reversible, we have also produced an optimal gate count for the state synthesis problem. Indeed, if we let $U = W^\dagger$, then we have $U|0\rangle = |\psi\rangle$. Moreover, if we label $p(n) = (d^n - 1)/(d - 1)$, then $U = \prod_{k=1}^{p(n)} \wedge(C(k), V(k)^\dagger)$ costs $p(n) = O(d^n)$ gates.

We postpone applications to §7 and next prove that Algorithm 3 is correct. The proof is new and is organized in terms of the \clubsuit -sequence.

5.4 Proof that \clubsuit -Householder achieves $W|\psi\rangle = |m\rangle$

For simplicity, we take $m = 0$, neglecting the permutations. Given n , $p(n) = (d^n - 1)/(d - 1)$ is the number of elements of the \clubsuit -sequence. It would suffice to prove (i) that each operator $\wedge[C(j), V(j)]$ guarantees $d - 1$ new zeroes in the state $|\psi_{j+1}\rangle$ not guaranteed in $|\psi_j\rangle$ and (ii) moreover that $\wedge[C(j), V(j)]$ does not act on previously guaranteed zeroes. The assertion (i) is straightforward and left to the reader; see Figure 2 caption. However, the second assertion is *false*. Rather, the controlled one-qudit operators do act on previously zeroed entries, but always replace them with a zero result. We next make this assertion precise and prove it.

Define the index set $S = \{0, 1, \dots, d^n - 1\}$ and introduce two new sets of dit-strings:

- $S_*(j)$ is the set of dit-strings for which the corresponding amplitude of $|\psi_j\rangle$ is not guaranteed zero by some $\wedge[C(k), V(k)]$, $k < j$.

- $S[C(j)]$ is the set of dit-strings that match $C(j)$, per Definition 5.1.

Also, define ℓ to be the index of the target symbol in $C(j)$: $C(j)_\ell = T$. Now there is a group action of $\mathbb{Z}/d\mathbb{Z}$ on the index set S corresponding to addition mod d on the ℓ^{th} dit:

$$c \bullet_\ell c_1 c_2 \dots c_n = c_1 c_2 \dots c_{\ell-1} (c_\ell + c \bmod d) c_{\ell+1} \dots c_n. \quad (20)$$

Since the operator $V(j)$ is applied to qudit ℓ , the amplitudes (components) of $|\psi_{j+1}\rangle$ are either equal to the corresponding amplitude of $|\psi_j\rangle$ or else are linear combinations of the $|\psi_j\rangle$ -amplitudes whose indices lie in the $\mathbb{Z}/d\mathbb{Z}$ orbit contained in $S[C(j)]$. To establish the correctness of **♣Householder**, we will prove the following Proposition.

Proposition 5.2 $|\psi_{j+1}\rangle$ has at least $d - 1$ more guaranteed zero amplitudes than $|\psi_j\rangle$.

Since **♣Householder** sets $j = 1, \dots, (d^n - 1)/(d - 1)$, this means that the final $|\psi_j\rangle$ has a single nonzero element corresponding to $|0\rangle$ and state synthesis has been achieved. We prove this result using three lemmas.

First we write $S_*(j)$ as the union of the three sets $R_1(j)$, $R_2(j)$, and $R_3(j)$ which we now define.

Definition 5.3 Suppose the j^{th} term of the **♣**-sequence is given by $c_1 c_2 \dots c_{\ell-1} \clubsuit \dots \clubsuit$. We have $C(j)$ the corresponding control word, with $C(j)_\ell = T$. Consider the following three sets, noting $R_1(j)$ may be empty.

$$\begin{aligned} R_1(j) &= \bigsqcup_{q=0}^{\ell-2} \left\{ c_1 c_2 \dots c_q k 00 \dots 0 ; k < c_{q+1}, k \in \{0, 1, \dots, d-1\} \right\} \\ R_2(j) &= \left\{ c_1 \dots c_{\ell-1} k 00 \dots 0 ; k \in \{0, 1, \dots, d-1\} \right\} \\ R_3(j) &= \left\{ f_1 \dots f_{\ell-1} k_\ell k_{\ell+1} \dots k_n ; f_1 f_2 \dots f_{\ell-1} > c_1 c_2 \dots c_{\ell-1}, k_* \in \{0, 1, \dots, d-1\} \right\} \end{aligned} \quad (21)$$

These sets may be interpreted in terms of Figure 2. Recall the figure recovers the **♣**-sequence by doing a depth-first search of the tree. In this context, $S_*(j)$ is the set of possibly nonzero components of $|\psi_j\rangle$ at the j^{th} node. The subset $R_3(j)$ results from indices that lie in nodes not yet traversed, loosely above the present node in the tree or to the right. The set $R_2(j)$ is precisely the set of indices in the current node, node j . The set $R_1(j)$ is the set of indices of elements that have been previously used to zero other elements and still might remain nonzero themselves; it is the set of indices of elements that were always at the top of nodes already traversed in the depth-first search. Thus, $R_1(j)$ is loosely a set of entries within nodes to the left and perhaps below node j .

The first lemma, along with the third, is used to show that the algorithm does not harm previously-introduced zeroes.

Lemma 5.4 Suppose the ℓ^{th} letter of $C(j)$ is the target symbol T , and label $\tilde{S}_*(j) = R_1(j) \sqcup R_2(j) \sqcup R_3(j)$. Then

$$(\mathbb{Z}/d\mathbb{Z}) \bullet_\ell \tilde{S}_*(j) \cap S[C(j)] \subseteq \tilde{S}_*(j) \cap S[C(j)]. \quad (22)$$

Proof: Due to the choice of a single control on a dit to the right of position ℓ in the appropriate term of the **♣**-sequence, $R_1(j) \cap S[C(j)] = \emptyset$. On the other hand, a direct computation verifies that $(\mathbb{Z}/d\mathbb{Z}) \bullet_\ell R_2(j) \subset R_2(j)$ and also that $R_2(j) \cap S[C(j)] = R_2(j)$.

Finally, we argue that $(\mathbb{Z}/d\mathbb{Z}) \bullet_\ell R_3(j) \subset R_3(j)$. However, the following partition is in general nontrivial:

$$R_3(j) = \{R_3(j) \cap S[C(j)]\} \sqcup \{R_3(j) \cap (S - S[C(j)])\}. \quad (23)$$

Should $C(j)$ admit no control, we are done. If not, let $m < \ell$ be the control qudit. Then

$$R_3(j) \cap S[C(j)] = \left\{ f_1 \dots f_{\ell-1} k_\ell k_{\ell+1} \dots k_n ; \mathbf{f}_m = \mathbf{c}_m, f_1 \dots f_{\ell-1} > c_1 c_2 \dots c_{\ell-1}, k_* \in \{0, 1, \dots, d-1\} \right\}. \quad (24)$$

Hence the $\mathbb{Z}/d\mathbb{Z}$ action respects the partition of Equation 23 as well. \square

The second lemma shows that the algorithm produces $d - 1$ newly guaranteed zeroes at each step.

Lemma 5.5 *Let $C(j)$, ℓ , and $S_*(j)$ be as above, with $C(j)$ resulting from $c_1c_2 \dots c_{\ell-1} \clubsuit \dots \clubsuit$ of the \clubsuit -sequence. Let $Z = \{c_1c_2 \dots c_{\ell-1}k00 \dots 0 ; k \in \{1, 2, \dots, d-1\} \cap \mathbb{Z}\}$ be the elements zeroed by $\wedge(C(j), V(j))$. Then $R_1(j) \sqcup R_2(j) \sqcup R_3(j) = R_1(j+1) \sqcup R_2(j+1) \sqcup R_3(j+1) \sqcup Z$.*

Proof: We break our argument into two cases based on the value of $c_{\ell-1}$.

Case $c_{\ell-1} < d - 1$: The $(j+1)^{\text{st}}$ term of the \clubsuit -sequence is given by $c_1c_2 \dots (c_{\ell-1} + 1)00 \dots 0\clubsuit$. Note that for leaves of the tree, the buffering sequence of zeroes is vacuous.

$$\begin{aligned} R_1(j+1) &= R_1(j) \sqcup R_2(j) - Z, \\ R_2(j+1) \sqcup R_3(j+1) &= R_3(j). \end{aligned} \tag{25}$$

Hence $R_1(j) \sqcup R_2(j) \sqcup R_3(j) = R_1(j+1) \sqcup R_2(j+1) \sqcup R_3(j+1) \sqcup Z$.

Case $c_{\ell-1} = d - 1$: Suppose instead the j^{th} \clubsuit -sequence term is $c_1c_2 \dots c_{\ell-2}(d-1)\clubsuit \dots \clubsuit$, so that the $(j+1)^{\text{st}}$ term is $c_1c_2 \dots c_{\ell-2}\clubsuit \clubsuit \dots \clubsuit$. We note that $\{c_0c_1 \dots c_{\ell-2}(d-1)0 \dots 0\} \in R_2(j) \cap R_2(j+1)$.^{*} Then

$$\begin{aligned} R_1(j) &= R_1(j+1) \sqcup R_2(j+1) - \{c_0c_1 \dots c_{\ell-2}(d-1)0 \dots 0\}, \\ R_2(j) &= Z \sqcup \{c_0c_1 \dots c_{\ell-2}(d-1)0 \dots 0\}, \\ R_3(j) &= R_3(j+1). \end{aligned} \tag{26}$$

From the first two, $R_1(j) \sqcup R_2(j) = R_1(j+1) \sqcup R_2(j+1) \sqcup Z$. Hence $R_1(j) \sqcup R_2(j) \sqcup R_3(j) = R_1(j+1) \sqcup R_2(j+1) \sqcup R_3(j+1) \sqcup Z$. \square

The third lemma shows that the set we considered in Lemma 5.4 is indeed the set of guaranteed zeros.

Lemma 5.6 $S_*(j) = R_1(j) \sqcup R_2(j) \sqcup R_3(j)$ is the set of guaranteed zero amplitudes (components) of a generic $|\psi_j\rangle$.

Proof: The proof is by induction. For $j = 1$, we have

$$R_1(1) = \emptyset, \quad R_2(1) = \{00 \dots 0*\}, \quad R_3(1) = \{c_1c_2 \dots c_{n-1}* ; \text{ some } c_j > 0\}. \tag{27}$$

Hence the entire index set $S = S_*(1) = R_1(1) \sqcup R_2(1) \sqcup R_3(1)$.

Hence, we suppose by way of induction that $S_*(j) = R_1(j) \sqcup R_2(j) \sqcup R_3(j)$ and attempt to prove the similar statement for $j+1$. Now $\wedge[C(j), V(j)]$ will add new zeroes to the amplitudes (components) with indices Z by Lemma 5.5. On the other hand, $\wedge[C(j), V(j)]$ will not destroy any zero amplitudes existing in $S_*(j)$ due to the induction hypothesis and Lemma 5.4. Thus $S_*(j+1) = R_1(j+1) \sqcup R_2(j+1) \sqcup R_3(j+1)$. \square

Proof of 5.2: The main result now follows after combining our three lemmas. \square

6 Unitary synthesis by reduction to triangular form

In this section, we present an asymptotically optimal unitary circuit not found in [4]. It leans heavily on the optimal state-synthesis of \clubsuit Householder. Since this state-synthesis circuit can likewise clear any length d^n vector using fewer than d^n single controls, the asymptotic is perhaps unsurprising. Yet the unitary circuit requires highly-controlled one-qudit unitary operators when clearing entries near the diagonal. Optimality persists since these are used sparingly. Two themes should be made clear at the outset:

- We process the size $d^n \times d^n$ unitary V in subblocks of size $d^{n-1} \times d^{n-1}$.
- Due to rank considerations, at least one block in each block-column of size $d^n \times d^{n-1}$ must remain full rank throughout.

^{*}So in the application, the amplitude (component) of this index is the single amplitude not zeroed by $\wedge[C(j), V(j)]$, but it is immediately afterwards zeroed by $\wedge[C(j+1), V(j+1)]$.

Hence, we cannot carelessly zero subcolumns. One solution is to triangularize the $d^{n-1} \times d^{n-1}$ matrices on the block diagonal, recursively. Given that strategy, the counts below show only $O(n^2 d^n)$ fully $(n-1)$ controlled one-qudit operations appear in the algorithm. This is allowed when working towards an asymptotic of $O(d^{2n})$ gates.

The organization for the algorithm is then as follows. Processing (triangularization) of V moves along block-columns of size $d^n \times d^{n-1}$ from left to right. In each block-column, we first triangularize the block $d^{n-1} \times d^{n-1}$ block-diagonal element, perhaps adding a control on the most significant qudit to a circuit produced by recursive triangularization. After this recursion, we zero the blocks below the block-diagonal element one column at a time. For each column j , $0 \leq j \leq d^{n-1} - 1$, the zeroing process is to collapse the $d^{n-1} \times 1$ subcolumns onto their j^{th} entries, again adding a control on the most significant qudit to prevent destroying earlier work. These subcolumn collapses produce the bulk of the zeroes and are done using **♣Householder**. After this, fewer than d entries remain to be zeroed in the column below the diagonal. These are eliminated using a controlled reflection containing $n-1$ controls and targeting the top line.

We now give a formal statement of the algorithm. We emphasize the addition of controls when previously generated circuits are incorporated into the universal circuit (i.e. recursively telescoping control.)

Algorithm 4: $\text{Triangle}(U, d, n)$

```

if  $n = 1$  then
    Triangularize  $U$  using a  $QR$  reduction.
else
    Reduce top-left  $d^{n-1} \times d^{n-1}$  subblock using Triangle  $(*, d, n-1)$ , (writing output to bottom
     $n-1$  circuit lines)
    for  $m = 0, 1, \dots, d-1$  do    % Block-column iteration
        for columns  $j = md^{n-1}, \dots, [(m+1)d^{n-1} - 1]$  do
            for  $\ell = (m+1), \dots, (d-1)$  do % Block-row iterate
                Use ♣Householder to zero the column entries  $(m+\ell)d^{n-1}, \dots, [(m+\ell+1)d^{n-1} - 1]$ ,
                leaving a nonzero entry at  $(m+\ell)c_2 \dots c_n$  for  $j = c_1 c_2 \dots c_n$  and
                adding  $|m+\ell\rangle$ - control on the most significant qudit.
            end for
            Clear the remaining nonzero entries below diagonal using one  $\wedge(Tc_2 \dots c_n, V)$ .
        end for    % All subdiagonal entries zero in block-col
        Use Triangle  $(*, d, n-1)$  on the  $d^{n-1} \times d^{n-1}$  matrix at the  $(m+1)^{\text{st}}$  block diagonal
        adding  $|m+1\rangle$ - control to the most significant qudit.
    end for
end if-else

```

To generate a circuit for a unitary operator U , we use **Triangle** to reduce U to a diagonal operator $W = \sum_{j=0}^{d^n-1} e^{i\phi_j} |j\rangle \langle j|$. Now V and $U = WV$ would be indistinguishable if a von Neumann measurement $\{|j\rangle \langle j|\}_{j=0}^{d^n-1}$ were made after each computation. However, the diagonal is important if U is a computation corresponding to a subblock of the circuit of a larger computation with other trailing, entangling interactions. In this case, the diagonal unitary can be simulated with $d^n \wedge_{n-1}(V)$ gates. Writing j in its d -ary expansion, $j = j_0 j_1 \dots j_{n-1}$ we have $W = \prod_{j=0}^{d^n-1} \otimes_{k=1}^n \text{INC}_k^{j_k} \wedge_{n-1} (e^{i\phi_j |d-1\rangle \langle d-1|}) \otimes_{k=1}^n \text{INC}_j^{-j_k}$. By the argument in §4, the gate count for such a simulation is $O[d^n(n-1)^{2+\log_2 d}]$. This is asymptotically irrelevant compared to the lower bound.

6.1 Counting gates and controls

Let $h(n, k)$ be the number of k -controls required in the **Single-♣Householder** reduction of some $|\psi\rangle \in \mathcal{H}(n, d)$. Then clearly $h(n, k) = 0$ for $k \geq 2$. Moreover, each 0-control results from an element of the **♣**-sequence of the form $00 \dots 0 \text{♣} \text{♣} \dots \text{♣}$, and there are n such sequences. Thus, since the number of elements of the **♣**-sequence is $(d^n - 1)/(d - 1)$, we see that

$$\begin{cases} h(n, 1) &= (d^n - 1)/(d - 1) - n \\ h(n, 0) &= n \end{cases} \quad (28)$$

We next count controls in the matrix algorithm **Triangle**. We break the count into two pieces: g for the work outside the main diagonal blocks and f for the total work.

Let $g(n, k)$ be the number of k -controls applied in operations in each column that zero the matrix below the block diagonal; this is the total work in the **for** j loops of **Triangle**. We use **Single-♣Householder** $d(d-1)d^{n-1}/2$ times since there are $d(d-1)/2$ blocks of size $d^{n-1} \times d^{n-1}$ below the block diagonal, and we add a single control to those counted in h . The last statement in the loop is executed $d^n - d^{n-1}$ times. Therefore, letting δ_j^k be the Kronecker delta, the counts are

$$g(n, k) = \delta_k^{n-1}(d^n - d^{n-1}) + \frac{1}{2}d(d-1)d^{n-1}h(n-1, k-1) \quad (29)$$

Supposing $n \geq 3$, then we see that

$$g(n, k) = \begin{cases} d^n - d^{n-1}, & k = n-1 \\ 0, & n-1 \leq k \leq 3 \\ \frac{1}{2}d^n(d^{n-1} - 1) - \frac{1}{2}d^n(d-1)(n-1), & k = 2 \\ \frac{1}{2}d^n(d-1)(n-1), & k = 1 \\ 0, & k = 0 \end{cases} \quad (30)$$

Finally, let $f(n, k)$ be the total number of k -controlled operations in the **Triangle** reduction, including the block diagonals. This work includes that counted in g , plus a recursive call to **Triangle** before the **for** m loop, plus $(d-1)$ calls within the k loop, for a total of

$$f(n, k) = g(n, k) + f(n-1, k) + (d-1)f(n-1, k-1), \quad (31)$$

with $f(n, 0) = 1$ and $f(1, k) = 0$ for $n, k > 0$.

Using the recursive relation of Equation 31 and the counts of Equation 30, we next argue that **Triangle** has no more than $O(d^{2n})$ controls. The following lemma is helpful.

Lemma 6.1 *For sufficiently large n , we have $f(n, k) \leq d^{2n-k+4}$.*

Proof: By inspection of Equation 30, we see that $g(n, k) \leq (1/2)d^{2n-k+2}$ for all k and n large. Now $f(n, 0) = 1$, which we take as an inductive hypothesis while supposing $f(n-1, \ell) \leq d^{2n-2-\ell+4} = d^{2n-\ell+2}$. Thus, using the recursion relation of Equation 31,

$$\begin{aligned} f(n, k) &\leq \frac{1}{2}d^{2n-k+2} + d^{2n-k+2} + (d-1)d^{2n-k+3} \\ &= d^{2n-k+4} \left(\frac{1}{2d^2} + \frac{1}{d^2} + 1 - \frac{1}{d} \right). \end{aligned} \quad (32)$$

Now since $d > 3/2$, we must have $\frac{1}{d} > \frac{3}{2d^2}$, whence an inductive proof of the result. \square

By the results from §4, each k -controlled single-qudit unitary operator costs $c_k = O[(k+2)^{2+\log_2(d)}]$ CINC and CINC^{-1} gates without ancillas. The expected number of CINC gates ℓ_T for the algorithm **Triangle** is then given by the weighted sum for the k -control gates in the diagonalization and the d^n instances of $n-1$ -controlled phase gates for emulation of the diagonal:

$$\begin{aligned} \ell_T &= d^n c_{n-1} + \sum_{k=0}^{n-1} c_k f(n, k) \\ &\leq 2(n+1)^{2+\log_2(d)} d^{n+4} + d^{8+2n} \sum_{k=0}^{n-1} d^{-k} k^{2+\lceil \log_2 d \rceil} \\ &\leq 2(n+1)^{2+\log_2(d)} d^{n+4} + d^{8+2n} \text{Li}_{-(2+\lceil \log_2 d \rceil)}(1/d) \\ &\leq 2(n+1)^{2+\log_2(d)} d^{n+4} + 26d^{8+2n}. \end{aligned} \quad (33)$$

In the third line we have used the fact that for the Polylogarithm function, $\text{Li}_{-(2+\lceil \log_2 d \rceil)}(1/d) \leq \text{Li}_{-3}(1/2) = 26$.

6.2 Comparison with the spectral algorithm

In an earlier work [4], we described an different algorithm for unitary synthesis. That algorithm relied on a spectral decomposition of the unitary and was also shown to be asymptotically optimal. For a circuit without ancillas, the CINC gate count ℓ_S using the spectral algorithm is:

$$\ell_S \leq 2d^{n+1}[(d^n - 1)/(d-1) - n] + (n+1)^{2+\log_2 d} d^{n+4} \quad (34)$$

In Table 2 the exact gate counts resulting from our implementations for unitary synthesis using **Triangle** and the spectral algorithm are tabulated. The result is that for a system with no ancillary resources, the spectral algorithm outperforms **Triangle** when the number of qudits n is greater than two. The general d^{2n} scaling for both is shown in Figure 3.

n	d	2	3	4	5	6	7	8	9	10
2		18 18	78 78	220 220	495 495	996 996	1 708 1 708	2 808 2 808	4 365 4 365	6 490 6 490
3		192 154	2 025 1 944	10 752 10 496	39 375 38 750	114 048 112 752	280 917 278 516	614 400 610 304	1 226 907 1 220 346	2 280 000 2 270 000
4		1 152 1 056	23 085 22 113	200 704 195 584	1 096 875 1 078 125	4 447 872 4 393 440	14 638 897 14 504 441	41 287 680 40 992 768	103 394 799 102 804 309	235 600 000 234 500 000
5		5 504 4 928	223 074 211 410	3 317 760 3 215 360	27 875 000 27 312 500	161 523 072 159 236 928	720 717 774 713 188 238	2 649 227 264 2 627 993 600	8 386 138 980 8 332 994 880	23 574 000 000 23 453 000 000
6		23 296 21 120	1 931 121 1 856 763	50 003 968 49 070 080						
7		92 672 84 224	16 605 891 16 087 572							
8		353 280 324 096	141 599 502 138 627 369							
9		1 333 248 1 246 208	1 224 144 819 1 209 914 010							
10		5 025 792 4 786 176	10 741 839 786 10 680 015 483							
11		19 128 320 18 452 480	95 432 986 134 95 147 070 876							
12		73 515 008 71 639 040								

Table 2: Exact gate counts for unitary synthesis without ancillas as a function of the number, n , and dimension, d , of the qudits. Each cell of the table lists the count for CINC and CINC⁻¹ gates using the most efficient of the two algorithms presented in the text. Boldface entries indicate that the **Triangle** algorithm was the most efficient, normal face type corresponds to counts using the spectral algorithm.

There are situations where **Triangle** may be preferred over the spectral algorithm. The later requires a classical diagonalization of the unitary U which requires $O(d^{3n})$ steps. For matrices of large size, particularly when there are degenerate eigenstates, numerical stability can be an issue. The classical computations involved in **Triangle** also scale like $O(d^{3n})$ but are carried out directly in the logical basis of the qudits.

7 Two applications of state synthesis

A primary motivation for describing state synthesis circuits is to utilize them as subcircuits for unitary synthesis as in §6. Yet there are also independent applications for the state-synth algorithm. We present two such.

7.1 Computing expected values

First, consider the problem of computing the expectation value of a Hermitian operator $A \in \mathcal{H}(n, d)$ i.e. $A \in \text{End}[\mathcal{H}(n, d)] \cong \mathbb{C}^{d^n \times d^n}$ with $A^\dagger = A$. For a system in the possibly mixed state ρ of n qudits, the the expectation of an operator A is $\langle A \rangle = \text{Tr}[A\rho]$. In some cases there does not exist a physically realistic direct measurement of A . However, one may infer the expectation value by a suitably weighted set of von Neumann measurements as follows. By the spectral theorem, any normal operator A may be diagonalized by a unitary transformation U : $A = U^\dagger D U$ where $D = \sum_{j=0}^{d^n-1} \lambda_j |j\rangle \langle j|$ and $\{\lambda_j\}_{j=0}^{d^n-1}$ are the eigenvalues of A . Then

$$\langle A \rangle = \text{Tr}[A\rho] = \text{Tr}[DU\rho U^\dagger] = \sum_{j=0}^{d^n-1} \lambda_j \text{Tr}[|j\rangle \langle j| U\rho U^\dagger]. \quad (35)$$




Figure 3: Performance comparison of the two algorithms for unitary synthesis on $n = 4$ qudits as a function of qudit dimension d . Triangles (boxes) indicate CINC gate counts for the **Triangle** (spectral) algorithm.

Hence we may compute $\langle A \rangle$ by performing three steps.

1. Prepare ρ .
2. Enact the unitary evolution U on ρ .
3. Perform the computational-basis von Neumann measurement on the resulting state, extracting all populations of the basis states $|j\rangle \langle j|$.

In some instances one may want to know the weight of a quantum state on a subspace of the operator A , i.e. $\langle P_S A P_S \rangle$ where P_S is some projection operator onto a subspace $\mathcal{H}_S \subseteq \mathcal{H}(n, d)$. In particular, consider the case of a k dimensional subspace diagonal in the eigenbasis $\{|u_j\rangle\}_{j=0}^{d^n-1}$ of A . We wish to compute $\text{Tr}[\sum_{j=1}^k \lambda_j |u_j\rangle \langle u_j| \rho]$ where $k < d^n$ and the eigenvalues of A have been reordered accordingly. Then we can rewrite the projection $P_S A P_S = \sum_{j=1}^k \lambda_j W(u_j) |j\rangle \langle j| W(u_j)^\dagger$ where $W(u_j)$ is a unitary extension of the mapping $|j\rangle \rightarrow |u_j\rangle$. The operator $W(u_j)$ is the unitary obtained in the state-synth algorithm. The expectation value is then

$$\langle P_S A P_S \rangle = \sum_{j=1}^k \lambda_j \text{Tr}[|j\rangle \langle j| W(u_j)^\dagger \rho W(u_j)]. \quad (36)$$

The expectation value can be measured as before but now one need only implement the state-synth operator k times on each state ρ of an ensemble of identically prepared states.

The above argument may in fact be generalized to compute the expectation value of any operator A . First decompose the operator as $A = A_h + A_a$ with $A_h = (A + A^\dagger)/2$ the Hermitian part and $A_a = (A - A^\dagger)/2$ the anti-Hermitian part of A . Both A_h and A_a are normal operators and therefore can be diagonalized. Hence, the expectation value can be computed by evaluating the weighted sum as per Eq. 35 and summing.

7.2 The general state synthesis problem

Both **Triangle** and the spectral algorithm are well adapted to the general state synthesis problem. This problem demands synthesizing any unitary extension of the many state mapping $\{|j\rangle \rightarrow |\psi_j\rangle \mid 0 \leq j \leq \ell \ll d^n\}$ [9]. It

is unclear what sorts of applications might arise when the states are arbitrary, requiring exponentially expensive circuits to build each. Nonetheless, less generic unitaries of this form have been used in quantum error correction to encode a few logical qudits into many physical qudits [8].

Triangle provides one solution to this problem. Start with a matrix containing $|\psi_j\rangle$ in its j th column, with “don’t care” entries in columns after column ℓ . Ignore any operations on the “don’t care” entries, and discard any gates meant to place zeros among them.

The spectral algorithm provides an alternative solution. Note that the matrix U formed from the product of the ℓ Householder transformations necessary to reduce the $d^n \times \ell$ matrix $[|\psi_1\rangle \dots |\psi_\ell\rangle]$ to diagonal form has $d^n - \ell$ eigenvalues equal to 1, so the spectral algorithm needs to build an eigenstate, apply a conditional phase to one logical basis ket, and unbuild the eigenstate only ℓ times.

8 Conclusions

This work concerns asymptotically optimal quantum circuits for qudits. By asymptotically optimal, we mean that the circuits require $O(d^n)$ gates of (no more than) two qudits for constructing arbitrary states and $O(d^{2n})$ gates for unitary evolutions. Contributions of this work are the following:

- We provide the first argument that both asymptotics survive even when no ancilla (helper) qudits are allowed.
- We present the state synthesis circuit in much more detail than previously published, in particular describing it in terms of iterates over a \clubsuit -sequence which plays a role similar to Gray codes for bits. Using the \clubsuit sequence, we provide the first proof that the state synthesis circuits actually achieve $U|0\rangle = |\psi\rangle$.
- We present **Triangle**, a new asymptotically optimal quantum circuit for qudit unitaries which is inspired by QR matrix factorization. Since it leans more heavily on QR than on spectral decomposition, the gate parameters of **Triangle** require less classical pre-processing than the spectral algorithm. Moreover, **Triangle** more closely resembles earlier quantum circuit design techniques [1, 14] than other asymptotically optimal qudit unitary circuits.
- §3 provides an elementary proof that $\{\text{CINC}\} \sqcup U(d)^{\otimes n}$ is exact-universal for qudits.

Some open questions remain. The $\wedge_1(V)$ gates are much better than earlier practice but not provably optimal, as is the case with qubits [12]. Moreover, the current best-practice n -qubit circuits exploit the cosine-sine decomposition (CSD), yet technical difficulties [13] with the tensor product structure make it quite unclear whether this matrix decomposition is useful for qudits.

Acknowledgements DPO received partial support from the National Science Foundation under Grants CCR-0204084 and CCF-0514213. GKB was supported in part by a grant from DARPA/QUIST. SSB was supported by a National Research Council postdoctoral fellowship.

References

- [1] A. Barenco, C. Bennett, R. Cleve, D. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter (1995), *Elementary Gates for Quantum Computation*, Phys. Rev. A, **52**, pp. 3457.
- [2] V. Bergholm, J. Vartiainen, M. Mtnen, and M. Salomaa (2005), Quantum circuits with uniformly controlled one-qubit gates, Phys. Rev. A, **71** pp. 052330.
- [3] G. Brennen, D. O’Leary, and S. Bullock (2005), *Criteria for exact qudit universality*, Phys. Rev. A, **71** pp. 052318.
- [4] S. Bullock, D. O’Leary, and G. Brennen (2005), *Asymptotically Optimal Quantum Circuits for d-level Systems*, Phys. Rev. Lett., **94**, pp. 230502.
- [5] G. Cybenko (2001), *Reducing Quantum Computations to Elementary Unitary Operations*, Comp. in Sci. and Eng., **27**, March/April.

- [6] D. Deutsch, A. Barenco, A. Ekert (1995), *Universality in Quantum Computation*, Proc. R. Soc. London A, **449**, pp. 669.
- [7] G.H. Golub and C. van Loan (1989), *Matrix Computations*, Johns Hopkins Press (Baltimore), 1989.
- [8] M. Grassl, M. Roetteler, T. Beth (2003), *Efficient Quantum Circuits for Non-Qubit Quantum Error-Correcting Codes*, International Journal of Foundations of Computer Science, **14**, pp. 757.
- [9] E. Knill (1995), *Approximation by Quantum Circuits*, <http://www.arxiv.org/abs/quant-ph/9508006>.
- [10] A. Muthukrishnan and C.R.Stroud Jr. (2000), *Multivalued Logic Gates for Quantum Computation*, Phys. Rev. A, **62**, pp. 052309.
- [11] V. Shende, S. Bullock, and I. Markov (2004), *Synthesis of quantum logic circuits*, IEEE Trans. on CAD, to appear, quant-ph/0406176.
- [12] G. Song and A. Klappenecker (2003), *Optimal realizations of controlled unitary gates*, Quantum Inf. Comput., **3(2)** pp. 139.
- [13] K.G.H. Vollbrecht and R.F. Werner (2000), *Why two qubits are special*, J. Math. Phys. **41**, pp. 6772.
- [14] J.J.Vartiainen, M.Möttönen, M.M.Salomaa (2004), *Efficient Decomposition of Quantum Gates*, Phys. Rev. Lett., **92**, pp. 177902.